

# **Financial Machine Learning**

**Other titles in Foundations and Trends® in Finance**

*Asset Allocation with Private Equity*

Arthur Korteweg and Mark M. Westerfield

ISBN: 978-1-68083-968-5

*Institutional Investors and Corporate Governance*

Amil Dasgupta, Vyacheslav Fos and Zacharias Sautner

ISBN: 978-1-68083-878-7

*The Implications of Heterogeneity and Inequality for Asset Pricing*

Stavros Panageas

ISBN: 978-1-68083-750-6

*Risk Sharing Within the Firm: A Primer*

Marco Pagano

ISBN: 978-1-68083-740-7

*The Economics of Credit Rating Agencies*

Francesco Sangiorgi and Chester Spatt

ISBN: 978-1-68083-380-5

*Initial Public Offerings: A Synthesis of the Literature and Directions for Future Research*

Michelle Lowry, Roni Michaely and Ekaterina Volkova

ISBN: 978-1-68083-340-9

# Financial Machine Learning

---

**Bryan Kelly**

Yale University and AQR Capital Management  
bryan.kelly@yale.edu

**Dacheng Xiu**

University of Chicago  
dacheng.xiu@chicagobooth.edu

**now**

the essence of knowledge

Boston — Delft

## Foundations and Trends<sup>®</sup> in Finance

*Published, sold and distributed by:*

now Publishers Inc.  
PO Box 1024  
Hanover, MA 02339  
United States  
Tel. +1-781-985-4510  
[www.nowpublishers.com](http://www.nowpublishers.com)  
[sales@nowpublishers.com](mailto:sales@nowpublishers.com)

*Outside North America:*

now Publishers Inc.  
PO Box 179  
2600 AD Delft  
The Netherlands  
Tel. +31-6-51115274

The preferred citation for this publication is

B. Kelly and D. Xiu. *Financial Machine Learning*. Foundations and Trends<sup>®</sup> in Finance, vol. 13, no. 3-4, pp. 205–363, 2023.

ISBN: 978-1-63828-291-4  
© 2023 B. Kelly and D. Xiu

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, mechanical, photocopying, recording or otherwise, without prior written permission of the publishers.

Photocopying. In the USA: This journal is registered at the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923. Authorization to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by now Publishers Inc for users registered with the Copyright Clearance Center (CCC). The 'services' for users can be found on the internet at: [www.copyright.com](http://www.copyright.com)

For those organizations that have been granted a photocopy license, a separate system of payment has been arranged. Authorization does not extend to other kinds of copying, such as that for general distribution, for advertising or promotional purposes, for creating new collective works, or for resale. In the rest of the world: Permission to photocopy must be obtained from the copyright owner. Please apply to now Publishers Inc., PO Box 1024, Hanover, MA 02339, USA; Tel. +1 781 871 0245; [www.nowpublishers.com](http://www.nowpublishers.com); [sales@nowpublishers.com](mailto:sales@nowpublishers.com)

now Publishers Inc. has an exclusive license to publish this material worldwide. Permission to use this content must be obtained from the copyright license holder. Please apply to now Publishers, PO Box 179, 2600 AD Delft, The Netherlands, [www.nowpublishers.com](http://www.nowpublishers.com); e-mail: [sales@nowpublishers.com](mailto:sales@nowpublishers.com)

**Foundations and Trends<sup>®</sup> in Finance**  
Volume 13, Issue 3-4, 2023  
**Editorial Board**

**Editor-in-Chief**

**Sheridan Titman**

University of Texas at Austin  
United States

**Associate Editors**

Josef Zechner

*WU Vienna University of Economics  
and Finance*

Chester Spatt

*Carnegie Mellon University*

## Editorial Scope

### Topics

Foundations and Trends® in Finance publishes survey and tutorial articles in the following topics:

- Corporate Finance
  - Corporate Governance
  - Corporate Financing
  - Dividend Policy and Capital Structure
  - Corporate Control
  - Investment Policy
  - Agency Theory and Information
- Financial Markets
  - Market Microstructure
  - Portfolio Theory
  - Financial Intermediation
  - Investment Banking
  - Market Efficiency
  - Security Issuance
  - Anomalies and Behavioral Finance
- Asset Pricing
  - Asset-Pricing Theory
  - Asset-Pricing Models
  - Tax Effects
  - Liquidity
  - Equity Risk Premium
  - Pricing Models and Volatility
  - Fixed Income Securities
- Derivatives
  - Computational Finance
  - Futures Markets and Hedging
  - Financial Engineering
  - Interest Rate Derivatives
  - Credit Derivatives
  - Financial Econometrics
  - Estimating Volatilities and Correlations

### Information for Librarians

Foundations and Trends® in Finance, 2023, Volume 13, 4 issues. ISSN paper version 1567-2395. ISSN online version 1567-2409. Also available as a combined paper and online subscription.

## Contents

---

<b>1</b>	<b>Introduction: The Case for Financial Machine Learning</b>	<b>2</b>
1.1	Prices are Predictions . . . . .	2
1.2	Information Sets are Large . . . . .	3
1.3	Functional Forms are Ambiguous . . . . .	5
1.4	Machine Learning versus Econometrics . . . . .	6
1.5	Challenges of Applying Machine Learning in Finance (and the Benefits of Economic Structure) . . . . .	8
1.6	Economic Content (Two Cultures of Financial Economics)	10
1.7	Roadmap . . . . .	15
<b>2</b>	<b>The Virtues of Complex Models</b>	<b>16</b>
2.1	Tools For Analyzing Machine Learning Models . . . . .	17
2.2	Bigger Is Often Better . . . . .	21
2.3	The Complexity Wedge . . . . .	27
<b>3</b>	<b>Return Prediction</b>	<b>30</b>
3.1	Data . . . . .	33
3.2	Experimental Design . . . . .	34
3.3	A Benchmark: Simple Linear Models . . . . .	39
3.4	Penalized Linear Models . . . . .	42
3.5	Dimension Reduction . . . . .	46
3.6	Decision Trees . . . . .	54

3.7	Vanilla Neural Networks . . . . .	61
3.8	Comparative Analyses . . . . .	65
3.9	More Sophisticated Neural Networks . . . . .	70
3.10	Return Prediction Models For “Alternative” Data . . . . .	72
<b>4</b>	<b>Risk-Return Tradeoffs</b>	<b>81</b>
4.1	APT Foundations . . . . .	81
4.2	Unconditional Factor Models . . . . .	82
4.3	Conditional Factor Models . . . . .	90
4.4	Complex Factor Models . . . . .	96
4.5	High-frequency Models . . . . .	98
4.6	Alphas . . . . .	100
<b>5</b>	<b>Optimal Portfolios</b>	<b>107</b>
5.1	“Plug-in” Portfolios . . . . .	109
5.2	Integrated Estimation and Optimization . . . . .	113
5.3	Maximum Sharpe Ratio Regression . . . . .	114
5.4	High Complexity MSRR . . . . .	118
5.5	SDF Estimation and Portfolio Choice . . . . .	120
5.6	Trading Costs and Reinforcement Learning . . . . .	128
<b>6</b>	<b>Conclusions</b>	<b>134</b>
	<b>Acknowledgements</b>	<b>136</b>
	<b>References</b>	<b>137</b>



# Financial Machine Learning

Bryan Kelly<sup>1</sup> and Dacheng Xiu<sup>2</sup>

<sup>1</sup>*Yale School of Management, AQR Capital Management, and NBER, USA; [bryan.kelly@yale.edu](mailto:bryan.kelly@yale.edu)*

<sup>2</sup>*University of Chicago Booth School of Business, USA; [dacheng.xiu@chicagobooth.edu](mailto:dacheng.xiu@chicagobooth.edu)*

---

## ABSTRACT

We survey the nascent literature on machine learning in the study of financial markets. We highlight the best examples of what this line of research has to offer and recommend promising directions for future research. This survey is designed for both financial economists interested in grasping machine learning tools, as well as for statisticians and machine learners seeking interesting financial contexts where advanced methods may be deployed.

---

# 1

---

## Introduction: The Case for Financial Machine Learning

---

### 1.1 Prices are Predictions

Modern analysis of financial markets centers on the following definition of a price, derived from the generic optimality condition of an investor:

$$P_{i,t} = \mathbb{E}[M_{t+1}X_{i,t+1}|\mathcal{I}_t]. \quad (1.1)$$

In words, the prevailing price of an asset,  $P_{i,t}$ , reflect investors' valuation of its future payoffs,  $X_{i,t+1}$ . These valuations are discounted based on investors' preferences, generically summarized as future realized marginal rates of substitution,  $M_{t+1}$ . The price is then determined by investor expectations of these objects given their conditioning information  $\mathcal{I}_t$ . In other words, prices are predictions—they reflect investors' best guesses for the (discounted) future payoffs shed by an asset.

It is common to analyze prices in an equivalent expected return, or “discount rate,” representation that normalizes (1.1) by the time  $t$  price:

$$\mathbb{E}[R_{i,t+1}|\mathcal{I}_t] = \beta_{i,t}\lambda_t, \quad (1.2)$$

where  $R_{i,t+1} = X_{i,t+1}/P_{i,t} - R_{f,t}$  is the asset's excess return,  $R_{f,t} = \mathbb{E}[M_{t+1}|\mathcal{I}_t]^{-1}$  is the one-period risk-free rate,  $\beta_{i,t} = \frac{\text{Cov}[M_{t+1}, R_{i,t+1}|\mathcal{I}_t]}{\text{Var}[M_{t+1}|\mathcal{I}_t]}$  is the asset's covariance with  $M_{t+1}$ , and  $\lambda_t = -\frac{\text{Var}[M_{t+1}|\mathcal{I}_t]}{\mathbb{E}[M_{t+1}|\mathcal{I}_t]}$  is the price

of risk. We can ask economic questions in terms of either prices or discount rates, but the literature typically opts for the discount rate representation for a few reasons. Prices are often non-stationary while discount rates are often stationary, so when the statistical properties of estimators rely on stationarity assumptions it is advantageous to work with discount rates. Also, uninteresting differences in the scale of assets' payoffs will lead to uninteresting scale differences in prices. But discount rates are typically unaffected by differences in payoff scale so the researcher need not adjust for them.

More generally, studying market phenomena in terms of returns alleviates some of the researcher's modeling burden by partially homogenizing data to have tractable dynamics and scaling properties. Besides, discount rates are also predictions, and their interpretation is especially simple and practically important.  $E[R_{i,t+1}|\mathcal{I}_t]$  describes investors' expectations for the appreciation in asset value over the next period. As such, the expected return is a critical input to allocation decisions. If we manage to isolate an empirical model for this expectation that closely fits the data, we have achieved a better understanding of market functionality and simultaneously derived a tool to improve resource allocations going forward. This is a fine example of duality in applied social science research: A good model both elevates scientific understanding and improves real-world decision-making.

## 1.2 Information Sets are Large

There are two conditions of finance research that make it fertile soil for machine learning methods: large conditioning information sets and ambiguous functional forms. Immediately evident from (1.1) is that the study of asset prices is inextricably tied to information. Guiding questions in the study of financial economics include "what information do market participants have and how do they use it?" The predictions embodied in prices are shaped by the available information that is pertinent to future asset payoffs ( $X_{i,t+1}$ ) and investors' preferences over those payoffs ( $M_{t+1}$ ). If prices behaved the same in all states of the world—e.g. if payoffs and preferences were close to i.i.d.—then information sets would drop out. But even the armchair investor dabbling

in their online account or reading the latest edition of *The Wall Street Journal* quickly intuits the vast scope of conditioning information lurking behind market prices. Meanwhile, the production function of the modern asset management industry is a testament to the vast amount of information flowing into asset prices: Professional managers (in various manual and automated fashions) routinely pore over troves of news feeds, data releases, and expert predictions in order to inform their investment decisions.

The expanse of price-relevant information is compounded by the panel nature of financial markets. The price of any given asset tends to vary over time in potentially interesting ways—this corresponds to the time series dimension of the panel. Meanwhile, at a given point in time, prices differ across assets in interesting ways—the cross section dimension of the panel. Time series variation in the market environment will affect many assets in interconnected ways. For example, most asset prices behave differently in high versus low risk conditions or in different policy regimes. As macroeconomic conditions change, asset prices adjust in unison through these common effects. Additionally, there are cross-sectional behaviors that are distinct to individual assets or groups of assets. So, conditioning information is not just time series in nature, but also includes asset-level attributes. A successful model of asset behavior must simultaneously account for shared dynamic effects as well as asset-specific effects (which may themselves be static or dynamic). As highlighted by Gu *et al.* (2020b),

*The profession has accumulated a staggering list of predictors that various researchers have argued possess forecasting power for returns. The number of stock-level predictive characteristics reported in the literature numbers in the hundreds and macroeconomic predictors of the aggregate market number in the dozens.*

Furthermore, given the tendency of financial economics research to investigate one or a few variables at a time, we have presumably left much ground uncovered. For example, only recently has the information content of news text emerged as an input to empirical models of (1.1), and there is much room for expansion on this frontier and others.

### 1.3 Functional Forms are Ambiguous

If asset prices are expectations of future outcomes, then the statistical tools to study prices are forecasting models. A traditional econometric approach to financial market research (e.g. Hansen and Singleton, 1982) first specifies a functional form for the return forecasting model motivated by a theoretical economic model, then estimates parameters to understand how candidate information sources associate with observed market prices within the confines of the chosen model. But which of the many economic models available in the literature should we impose?

The formulation of the first-order condition, or “Euler equation,” in (1.1) is broad enough to encompass a wide variety of structural economic assumptions. This generality is warranted because there is no consensus about which specific structural formulations are viable. Early consumption-based models fail to match market price data by most measures (e.g. Mehra and Prescott, 1985). Modern structural models match price data somewhat better if the measure of success is sufficiently forgiving (e.g. Chen *et al.*, 2022a), but the scope of phenomena they describe tends to be limited to a few assets and is typically evaluated only on an in-sample basis.

Given the limited empirical success of structural models, most empirical work in the last two decades has opted away from structural assumptions to less rigid “reduced-form” or “no-arbitrage” frameworks. While empirical research of markets often steers clear of imposing detailed economic structure, it typically imposes statistical structure (for example, in the form of low-dimensional factor models or other parametric assumptions). But there are many potential choices for statistical structure in reduced-form models, and it is worth exploring the benefits of flexible models that can accommodate many different functional forms and varying degrees of nonlinearity and variable interactions.

Enter machine learning tools such as kernel methods, penalized likelihood estimators, decision trees, and neural networks. Comprised of diverse nonparametric estimators and large parametric models, machine learning methods are explicitly designed to approximate unknown data generating functions. In addition, machine learning can help integrate many data sources into a single model. In light of the discussion in

Section 1.2, effective modeling of prices and expected returns requires rich conditioning information in  $\mathcal{I}_t$ . On this point, Cochrane (2009)<sup>1</sup> notes that “*We obviously don’t even observe all the conditioning information used by economic agents, and we can’t include even a fraction of observed conditioning information in our models.*” Hansen and Richard (1987) (and more recently Martin and Nagel, 2021) highlight differences in information accessible to investors inside an economic model versus information available to an econometrician on the outside of a model looking in. Machine learning is a toolkit that can help narrow the gap between information sets of researchers and market participants by providing methods that allow the researcher to assimilate larger information sets.

The more expansive we can be in our consideration of large conditioning sets, the more realistic our models will be. This same logic applies to the question of functional form. Not only do market participants impound rich information into their forecasts, they do it in potentially complex ways that leverage the nuanced powers of human reasoning and intuition. We must recognize that investors use information in ways that we as researchers cannot know explicitly and thus cannot exhaustively (and certainly not concisely) specify in a parametric statistical model. Just as Cochrane (2009) reminds us to be circumspect in our consideration of conditioning information, we must be equally circumspect in our consideration of functional forms.

## 1.4 Machine Learning versus Econometrics

What is machine learning, and how is it different from traditional econometrics? Gu *et al.* (2020b) emphasize that the definition of machine learning is inchoate and the term is at times corrupted by the marketing purposes of the user. We follow Gu *et al.* (2020b) and use the term to describe (i) a diverse collection of high-dimensional models for statistical prediction, combined with (ii) “regularization” methods for model selection and mitigation of overfit, and (iii) efficient algorithms for searching among a vast number of potential model specifications.

---

<sup>1</sup>Readers of this survey are encouraged to re-visit chapter 8 of Cochrane (2009) and recognize the many ways machine learning concepts mesh with his outline of the role of conditioning information in asset prices.

Given this definition, it should be clear that, in any of its incarnations, financial machine learning amounts to a set of procedures for estimating a statistical model and using that model to make decisions. So, at its core, machine learning *need not be differentiated* from econometrics or statistics more generally. Many of the ideas underlying machine learning have lived comfortably under the umbrella of statistics for decades (Israel *et al.*, 2020).

In order to learn through the experience of data, the machine needs a functional representation of what it is trying to learn. The researcher must make a representation choice—this is a canvas upon which the data will paint its story. Part (i) of our definition points out that machine learning brings an open-mindedness to functional representations that are highly parameterized and often nonlinear. Small models are rigid and oversimplified, but their parsimony has benefits like comparatively precise parameter estimates and ease of interpretation. Large and sophisticated models are much more flexible, but can also be more sensitive and suffer from poor out-of-sample performance when they overfit noise in the system. Researchers turn to large models when they believe the benefits from more accurately describing the complexities of real world phenomena outweigh the costs of potential overfit. At an intuitive level, machine learning is a way to pursue statistical analysis when the analyst is unsure which specific structure their statistical model should take. In this sense, much of machine learning can be viewed as nonparametric (or semi-parametric) modeling. Its *modus operandi* considers a variety of potential model specifications and asks the data's guidance in choosing which model is most effective for the problem at hand. One may ask: when does the analyst *ever* know what structure is appropriate for their statistical analysis? The answer of course is “never,” which is why machine learning is generally valuable in financial research. As emphasized by Breiman (2001), its focus on maximizing prediction accuracy in the face of an unknown data model is the central differentiating feature of machine learning from the traditional statistical objective of estimating a known data generating model and conducting hypothesis tests.

Part (ii) of our definition highlights that machine learning chooses a preferred model (or combination of models) from a “diverse collection”

of candidate models. Again, this idea has a rich history in econometrics under the heading of model selection (and, relatedly, model averaging). The difference is that machine learning puts model selection at the heart of the empirical design. The process of searching through many models to find top performers (often referred to as model “tuning”) is characteristic of all machine learning methods. Of course, selecting from multiple models mechanically leads to in-sample overfitting and can produce poor out-of-sample performance. Thus machine learning research processes are accompanied by “regularization,” which is a blanket term for constraining model size to encourage stable performance out-of-sample. As Gu *et al.* (2020b) put it, “*An optimal model is a ‘Goldilocks’ model. It is large enough that it can reliably detect potentially complex predictive relationships in the data, but not so flexible that it is dominated by overfit and suffers out-of-sample.*” Regularization methods encourage smaller models; richer models are only selected if they are likely to give a genuine boost to out-of-sample prediction accuracy.

Element (iii) in the machine learning definition is perhaps its clearest differentiator from traditional statistics, but also perhaps the least economically interesting. When data sets are large and/or models are very heavily parameterized, computation can become a bottleneck. Machine learning has developed a variety of approximate optimization routines to reduce computing loads. For example, traditional econometric estimators typically use all data points in every step of an iterative optimization routine and only cease the parameter search when the routine converges. Shortcuts such as using subsets of data and halting a search before convergence often reduce computation and do so with little loss of accuracy (see, e.g., stochastic gradient descent and early stopping which are two staples in neural network training).

## 1.5 Challenges of Applying Machine Learning in Finance (and the Benefits of Economic Structure)

While financial research is in many ways ideally suited to machine learning methods, some aspects of finance also present challenges for machine learning. Understanding these obstacles is important for developing realistic expectations about the benefits of financial machine learning.



First, while machine learning is often viewed as a “big data” tool, many foundational questions in finance are frustrated by the decidedly “small data” reality of economic time series. Standard data sets in macro finance, for example, are confined to a few hundred monthly observations. This kind of data scarcity is unusual in other machine learning domains where researchers often have, for all intents and purposes, unlimited data (or the ability to generate new data as needed). In time series research, new data accrues only through the passage of time.

Second, financial research often faces weak signal-to-noise ratios. Nowhere is this more evident than in return prediction, where the forces of market efficiency (profit maximization and competition) are ever striving to eliminate the predictability of price movements (Samuelson, 1965; Fama, 1970). As a result, price variation is expected to emanate predominantly from the arrival of unanticipated news (which is unforecastable noise from the perspective of the model). Markets may also exhibit inefficiencies and investor preferences may give rise to time-varying risk premia, which result in some predictability of returns. Nonetheless, we should expect return predictability to be small and fiercely competed over.

Third, investors learn and markets evolve. This creates a moving target for machine learning prediction models. Previously reliable predictive patterns may be arbitrated away. Regulatory and technological changes alter the structure of the economy. Structural instability makes finance an especially complex learning domain and compounds the challenges of small data and low signal-to-noise ratios.

These challenges present an opportunity to benefit from knowledge gained by economic theory. As noted by Israel *et al.* (2020),

*“A basic principle of statistical analysis is that theory and model parameters are substitutes. The more structure you can impose in your model, the fewer parameters you need to estimate and the more efficiently your model can use available data points to cut through noise. That is, models are helpful because they filter out noise. But an over-simplified model can filter out some signal too, so in a data-rich and high signal-to-noise environment, you would not want to use*

*an unnecessarily small model. One can begin to tackle small data and low signal-to-noise problems by bringing economic theory to describe some aspects of the data, complemented by machine learning tools to capture aspects of the data for which theory is silent.”*

Economic theory can be fused with machine learning by imposing theory-implied cross-parameter restrictions in the statistical model specification. For example, a minimal yet potentially powerful theoretical restriction to impose on a machine learning model is the absence of arbitrage (e.g. Cao *et al.*, 2021). Another example of a theoretical restriction that may be worth imposing is that only systematic risk is compensated. In a machine learning factor model, for example, this can be achieved by cross-parameter restrictions that anchor assets’ mean returns to their factor betas (while still allowing betas to have a flexible machine learning functional form, as in Gu *et al.*, 2020a).

## 1.6 Economic Content (Two Cultures of Financial Economics)

We recall Breiman (2001)’s essay on the “two cultures” of statistics, which has an analogue in financial economics (with appropriate modifications). One is the “structural model/hypothesis test” culture, which favors imposing fully or partially specified structural assumptions and investigating economic mechanisms through hypothesis tests. The traditional program of empirical asset pricing analysis (pre-dating the emergence of reduced form factor models and machine learning prediction models) studies prices through the lens of heavily constrained prediction models. The constraints come in the form of i) specific functional forms/distributions, and ii) limited variables admitted into the conditioning information set. These models often “generalize” poorly in the sense that they have weak explanatory power for asset price behaviors outside the narrow purview of the model design or beyond the training data set. This is such an obvious statement that one rarely considers out-of-sample performance of fully specified structural asset pricing models.

The other is the “prediction model” culture, which values statistical explanatory power above all else, and is born largely from the limitations of the earlier established structural culture. The prediction model culture willingly espouses model specifications that might lack an explicit association with economic theory, so long as they produce meaningful, robust improvements in data fit versus the status quo.<sup>2</sup> In addition to reduced-form modeling that has mostly dominated empirical finance since the 1990’s, financial machine learning research to date falls squarely in this second culture.

“There is no economics” is a charge sometimes lobbed at the statistical prediction research by economic seminar audiences, discussants, and referees. This criticism is often misguided and we should guard against it unduly devaluing advancements in financial machine learning. Let us not miss the important economic role of even the purest statistical modeling applications in finance. Relatively unstructured prediction models makes them no less economically important than the traditional econometrics of structural hypothesis testing, they just play a different scientific role. Hypothesis testing learns economics by probing specific economic mechanisms. But economics is not just about testing theoretical mechanisms. Atheoretical (for lack of a better term) prediction models survey the empirical landscape in broader terms, charting out new empirical facts upon which theories can be developed, and for which future hypothesis tests can investigate mechanisms. These two forms of empirical investigation—precision testing and general cartography—play complementary roles in the Kuhnian process of scientific advancement.

Consider the fundamental question of asset pricing research: What determines asset risk premia? Even if we could observe expected returns perfectly, we would still need theories to explain their behavior and empirical analysis to test those theories. But we can’t observe risk premia, and they are stubbornly hard to estimate. Machine learning

---

<sup>2</sup>It remains critical to determine whether a candidate forecasting model genuinely improves predictive accuracy with reasonable confidence or is just the result of data mining. Likewise, the difficulty of “auditing” the model development process or tracking the number of iterations a researcher attempted to fine-tune results calls for continued vigilance for questionable research designs, as is also the case in traditional empirical finance research.

makes progress on measuring risk premia, which facilitates development of better theories of economic mechanisms that determine their behavior.

A critical benefit of expanding the set of known contours in the empirical landscape is that, even if details of the economic mechanisms remain shrouded, economic actors—financial market participants in particular—can always benefit from improved empirical maps. The prediction model culture has a long tradition of producing research to help investors, consumers, and policymakers make better decisions. Improved predictions provide more accurate descriptions of the state-dependent distributions faced by these economic actors.

Economics is by and large an applied field. The economics of the prediction model culture *lies precisely in* its ability to improve predictions. Armed with better predictions—i.e., more accurate assessments of the economic opportunity set—agents can better trade off costs and benefits when allocating scarce resources. This enhances welfare. Nowhere is this more immediately clear than in the portfolio choice problem. We may not always understand the economic mechanisms by which a model delivers better return or risk forecasts; but if it does, it boosts the utility of investors and is thus economically important.

Breiman's (2001) central criticism of the structural hypothesis test culture is that:

*“when a model is fit to data to draw quantitative conclusions: the conclusions are about the model’s mechanism, and not about nature’s mechanism. If the model is a poor emulation of nature, the conclusions may be wrong.”*

We view this less as a criticism of structural modeling, which must remain a foundation of empirical finance, but rather as a motivation and defense of prediction models. The two-culture dichotomy is, of course, a caricature. Research spans a spectrum and draws on multiple tools, and researchers do not separate into homogenous ideological camps. Both cultures are economically important. Breiman (2001) encourages us to consider flexible, even nonparametric, models to learn about economic mechanisms:

*“The point of a model is to get useful information about the relation between the response and predictor variables.*

*Interpretability is a way of getting information. But a model does not have to be simple to provide reliable information about the relation between predictor and response variables; neither does it have to be a [structural] data model.”*

Prediction models are a first step to understanding mechanisms. Moreover, structural modeling can benefit directly from machine learning without sacrificing pointed hypothesis tests or its specificity of economic mechanisms.<sup>3</sup> Thus far machine learning has predominantly served the prediction model culture of financial economics. It is important to recognize it as a similarly potent tool for the structural hypothesis testing culture (this is a critical direction for future machine learning research in finance). Surely, a research program founded solely on “measurement without theory” (Koopmans, 1947) is better served by also considering data through the lens of economic theory and with a deep understanding of the Lucas Jr (1976) critique. Likewise, a program that only interprets data through extant economic models can overlook unexpected yet economically important statistical patterns. And on the margin, machine learning models that are more parsimonious, transparent, and economically interpretable are also more desirable, just as in traditional statistical modeling.

Hayek (1945) confronts the economic implications of dispersed information for resource allocation. Regarding his central question of how to achieve an effective economic order, he notes:

*If we possess all the relevant information, if we can start out from a given system of preferences, and if we command complete knowledge of available means, the problem which remains is purely one of logic... This, however, is emphatically not the economic problem which society faces. And the economic calculus which we have developed to solve this logical problem, though an important step toward the solution of the economic problem of society, does not yet provide an answer to it. The reason for this is that the ‘data’ from which the economic calculus starts are never for the whole*

---

<sup>3</sup>See, for example, our discussion of Chen and Ludvigson (2009), in Section 5.5.

*society 'given' to a single mind which could work out the implications and can never be so given.*

While Hayek's main interest is in the merits of decentralized planning, his statements also have implications for information technologies in general, and prediction technologies in particular. Let us be so presumptuous as to reinterpret Hayek's statement as a statistical problem: There is a wedge between the efficiency of allocations achievable by economic agents when the data generating process (DGP) is known, versus when it must be estimated. First, there is the problem of model specification—economic agents simply cannot be expected to correctly specify their statistical models. They must use some form of mis-specified parametric model or a nonparametric approximating model. In either case, mis-specification introduces a wedge between the optimal allocations achievable when the DGP is known (call this “first-best”) and the allocations derived from their mis-specified models (call this “second-best”). But even second best is implausible, because we must estimate these models with finite data. This gives rise to yet another wedge, that due to sampling variation. Even if we knew the functional form of the DGP, we still must estimate it and noise in our estimates produces deviations from first-best. Compound that with the realism of mis-specification, and we recognize that in reality we must always live with “third-best” allocations; i.e., mis-specified models that are noisily estimated.

Improved predictions derived from methods that can digest vast information and data sets provide an opportunity to mitigate the wedges between the pure “logic” problem of first-best resource allocation noted by Hayek, and third-best realistic allocations achievable by economic agents. The wedges never shrink to zero due to statistical limits to learnability (Da *et al.*, 2022; Didisheim *et al.*, 2023). But powerful approximating models and clever regularization devices mean that machine learning is economically important exactly because it can lead to better decisions. The problem of portfolio choice is an illustrative example. A mean-variance investor who knows the true expected return and covariance matrix of assets simply executes the “logic” of a Markowitz portfolio and achieves a first-best allocation. But, in analogy to Hayek, this is emphatically *not* the problem that real world investors

grapple with. Instead, their problem is primarily one of measurement—one of prediction. The investor seeks a sensible expected return and covariance estimate that, when combined with the Markowitz objective, performs reasonably well out-of-sample. Lacking high-quality measurements, the Markowitz solution can behave disastrously, as much research has demonstrated.

## 1.7 Roadmap

This survey is organized as follows. Section 2 analyzes the theoretical benefits of highly parameterized machine learning models in financial economics. Section 3 surveys the variety of machine learning methods employed in the empirical analysis of asset return predictability. Section 4 focuses on machine learning analyses of factor pricing models and the resulting empirical conclusions for risk-return tradeoffs. Section 5 presents the role of machine learning in identifying optimal portfolios and stochastic discount factors. Section 6 offers brief conclusions and directions for future work.

## References

---

- Ahn, S. C. and J. Bae. (2022). “Forecasting with Partial Least Squares When a Large Number of Predictors are Available”. *Tech. rep.* Arizona State University and University of Glasgow.
- Ai, C. and X. Chen. (2003). “Efficient Estimation of Models with Conditional Moment Restrictions Containing Unknown Functions”. *Econometrica*. 71(6): 1795–1843.
- Ai, C. and X. Chen. (2007). “Estimation of possibly misspecified semi-parametric conditional moment restriction models with different conditioning variables”. *Journal of Econometrics*. 141(1): 5–43.
- Ait-Sahalia, Y. and M. W. Brandt. (2001). “Variable Selection for Portfolio Choice”. *The Journal of Finance*. 56: 1297–1351.
- Ait-Sahalia, Y., J. Fan, L. Xue, and Y. Zhou. (2022). “How and when are high-frequency stock returns predictable?” *Tech. rep.* Princeton University.
- Ait-Sahalia, Y., J. Jacod, and D. Xiu. (2021). “Continuous-Time Fama-MacBeth Regressions”. *Tech. rep.* Princeton University and the University of Chicago.
- Ait-Sahalia, Y., I. Kalnina, and D. Xiu. (2020). “High Frequency Factor Models and Regressions”. *Journal of Econometrics*. 216: 86–105.
- Ait-Sahalia, Y. and A. W. Lo. (1998). “Nonparametric estimation of state-price densities implicit in financial asset prices”. *The journal of finance*. 53(2): 499–547.



- Ait-Sahalia, Y. and A. W. Lo. (2000). “Nonparametric risk management and implied risk aversion”. *Journal of econometrics*. 94(1-2): 9–51.
- Ait-Sahalia, Y. and D. Xiu. (2017). “Using Principal Component Analysis to Estimate a High Dimensional Factor Model with High-Frequency Data”. *Journal of Econometrics*. 201: 388–399.
- Ait-Sahalia, Y. and D. Xiu. (2019). “Principal Component Analysis of High Frequency Data”. *Journal of the American Statistical Association*. 114: 287–303.
- Allen-Zhu, Z., Y. Li, and Z. Song. (2019). “A convergence theory for deep learning via over-parameterization”. In: *International Conference on Machine Learning*. PMLR. 242–252.
- Altman, E. I. (1968). “Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy”. *The Journal of Finance*. 23(4): 589–609.
- Anders, U., O. Korn, and C. Schmitt. (1998). “Improving the pricing of options: A neural network approach”. *Journal of forecasting*. 17(5-6): 369–388.
- Andersen, T. G. and T. Bollerslev. (1998). “Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts”. *International Economic Review*. 39: 885–905.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys. (2001). “The Distribution of Exchange Rate Realized Volatility”. *Journal of the American Statistical Association*. 96: 42–55.
- Antoine, B., K. Proulx, and E. Renault. (2018). “Pseudo-True SDFs in Conditional Asset Pricing Models\*”. *Journal of Financial Econometrics*. 18(4): 656–714. DOI: [10.1093/jjfinec/nby017](https://doi.org/10.1093/jjfinec/nby017).
- Ao, M., L. Yingying, and X. Zheng. (2018). “Approaching Mean-Variance Efficiency for Large Portfolios”. *The Review of Financial Studies*. 32(7): 2890–2919.
- Arlot, S. and A. Celisse. (2010). “A survey of cross-validation procedures for model selection”. *Statistics surveys*. 4: 40–79.
- Aubry, M., R. Kraussl, M. Gustavo, and C. Spaenjers. (2023). “Biased Auctioneers”. *The Journal of Finance*. 78(2): 795–833.
- Avramov, D., S. Cheng, and L. Metzker. (2023a). “Machine Learning vs. Economic Restrictions: Evidence from Stock Return Predictability”. *Management Science*. 69(5): 2547–3155.

- Avramov, D., S. Cheng, L. Metzker, and S. Voigt. (2023b). “Integrating factor models”. *Journal of Finance*. 78(3): 1593–1646.
- Avramov, D., G. Kaplanski, and A. Subrahmanyam. (2022). “Postfundamentals Price Drift in Capital Markets: A Regression Regularization Perspective”. *Management Science*. 68(10): 7658–7681.
- Avramov, D. and G. Zhou. (2010). “Bayesian Portfolio Analysis”. *Annual Review of Financial Economics*. 2(1): 25–47.
- Bai, J. (2003). “Inferential theory for factor models of large dimensions”. *Econometrica*. 71(1): 135–171.
- Bai, J. and S. Ng. (2002). “Determining the Number of Factors in Approximate Factor Models”. *Econometrica*. 70: 191–221.
- Bai, J. and S. Ng. (2021). “Approximate Factor Models with Weaker Loading”. *Tech. rep.* Columbia University.
- Bajgrowicz, P. and O. Scaillet. (2012). “Technical trading revisited: False discoveries, persistence tests, and transaction costs”. *Journal of Financial Economics*. 106(3): 473–491.
- Baker, M. and J. Wurgler. (2006). “Investor sentiment and the cross-section of stock returns”. *The journal of Finance*. 61(4): 1645–1680.
- Baker, M. and J. Wurgler. (2007). “Investor Sentiment in the Stock Market”. *Journal of Economic Perspectives*. 21(2): 129–152.
- Balduzzi, P. and A. W. Lynch. (1999). “Transaction costs and predictability: some utility cost calculations”. *Journal of Financial Economics*. 52: 47–78.
- Bali, T. G., A. Goyal, D. Huang, F. Jiang, and Q. Wen. (2020). “Predicting Corporate Bond Returns: Merton Meets Machine Learning”. *Tech. rep.* Georgetown University.
- Bansal, R. and S. Viswanathan. (1993). “No Arbitrage and Arbitrage Pricing: A New Approach”. *The Journal of Finance*. 48(4): 1231–1262.
- Bansal, R. and A. Yaron. (2004). “Risks for the long run: A potential resolution of asset pricing puzzles”. *The journal of Finance*. 59(4): 1481–1509.
- Bao, W., J. Yue, and Y. Rao. (2017). “A deep learning framework for financial time series using stacked autoencoders and long-short term memory”. *PLOS ONE*. 12(7): 1–24.

- Barberis, N. (2018). “Psychology-based models of asset prices and trading volume”. In: *Handbook of behavioral economics: applications and foundations 1*. Vol. 1. Elsevier. 79–175.
- Barberis, N. and R. Thaler. (2003). “A survey of behavioral finance”. *Handbook of the Economics of Finance*. 1: 1053–1128.
- Barndorff-Nielsen, O. E. and N. Shephard. (2002). “Econometric Analysis of Realized Volatility and Its Use in Estimating Stochastic Volatility Models”. *Journal of the Royal Statistical Society, B*. 64: 253–280.
- Barras, L., O. Scaillet, and R. Wermers. (2010). “False discoveries in mutual fund performance: Measuring luck in estimated alphas”. *Journal of Finance*. 65(1): 179–216.
- Bartlett, P. L., P. M. Long, G. Lugosi, and A. Tsigler. (2020). “Benign overfitting in linear regression”. *Proceedings of the National Academy of Sciences*. 117(48): 30063–30070.
- Basu, S. (1977). “Investment Performance of Common Stocks in Relation to Their Price-Earnings Ratios: A Test of the Efficient Market Hypothesis”. *The Journal of Finance*. 32(3): 663–682.
- Belkin, M., D. Hsu, S. Ma, and S. Mandal. (2018). “Reconciling modern machine learning and the biasvariance trade-off. arXiv e-prints”.
- Belkin, M. (2021). “Fit without fear: remarkable mathematical phenomena of deep learning through the prism of interpolation”. *Acta Numerica*. 30: 203–248.
- Belkin, M., D. Hsu, and J. Xu. (2020). “Two models of double descent for weak features”. *SIAM Journal on Mathematics of Data Science*. 2(4): 1167–1180.
- Belkin, M., A. Rakhlin, and A. B. Tsybakov. (2019). “Does data interpolation contradict statistical optimality?” In: *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR. 1611–1619.
- Benjamini, Y. and Y. Hochberg. (1995). “Controlling the false discovery rate: a practical and powerful approach to multiple testing”. *Journal of the Royal Statistical Society. Series B (Methodological)*. 57(1): 289–300.

- Bianchi, D., M. Büchner, and A. Tamoni. (2021). “Bond risk premiums with machine learning”. *The Review of Financial Studies*. 34(2): 1046–1089.
- Bickel, P. J. and E. Levina. (2008a). “Covariance Regularization by Thresholding”. *Annals of Statistics*. 36(6): 2577–2604.
- Bickel, P. J. and E. Levina. (2008b). “Regularized Estimation of Large Covariance Matrices”. *Annals of Statistics*. 36: 199–227.
- Black, F. and R. Litterman. (1992). “Global Portfolio Optimization”. *Financial Analysts Journal*. 48(5): 28–43.
- Bollerslev, T., S. Z. Li, and V. Todorov. (2016). “Roughing up beta: Continuous versus discontinuous betas and the cross section of expected stock returns”. *Journal of Financial Economics*. 120: 464–490.
- Bollerslev, T., M. C. Medeiros, A. Patton, and R. Quaadvlieg. (2022). “From Zero to Hero: Realized Partial (Co)variances”. *Journal of Econometrics*. 231: 348–360.
- Bollerslev, T. and V. Todorov. (2011). “Estimation of jump tails”. *Econometrica*. 79(6): 1727–1783.
- Box, G. E. P., G. M. Jenkins, G. C. Reinsel, and G. M. Ljung. (2015). *Time Series Analysis: Forecasting and Control*. 5th. Wiley.
- Box, G. E. and G. Jenkins. (1970). *Time Series Analysis: Forecasting and Control*. San Francisco: Holden-Day.
- Brandt, M. W. (1999). “Estimating Portfolio and Consumption Choice: A Conditional Euler Equations Approach”. *The Journal of Finance*. 54(5): 1609–1645.
- Brandt, M. W. (2010). “Portfolio Choice Problems”. In: *Handbook of Financial Econometrics*. Ed. by Y. Ait-Sahalia and L. P. Hansen. Amsterdam, The Netherlands: North-Holland. 269–336.
- Brandt, M. W. and P. Santa-Clara. (2006). “Dynamic Portfolio Selection by Augmenting the Asset Space”. *The Journal of Finance*. 61(5): 2187–2217.
- Brandt, M. W., P. Santa-Clara, and R. Valkanov. (2009). “Covariance regularization by parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns”. *Review of Financial Studies*. 22: 3411–3447.

- Breiman, L. (1995). “The Mathematics of Generalization”. In: CRC Press. Chap. Reflections After Refereeing Papers for NIPS. 11–15.
- Breiman, L. (2001). “Random forests”. *Machine learning*. 45(1): 5–32.
- Britten-Jones, M. (1999). “The Sampling Error in Estimates of Mean-Variance Efficient Portfolio Weights”. *The Journal of Finance*. 54(2): 655–671.
- Brown, T., B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei. (2020). “Language Models are Few-Shot Learners”. In: *Advances in Neural Information Processing Systems*. Ed. by H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin. Vol. 33. Curran Associates, Inc. 1877–1901.
- Bryzgalova, S., V. DeMiguel, S. Li, and M. Pelger. (2023). “Asset-Pricing Factors with Economic Targets”. Available at SSRN 4344837.
- Bryzgalova, S., M. Pelger, and J. Zhu. (2020). “Forest through the Trees: Building Cross-Sections of Asset Returns”. *Tech. rep.* London School of Business and Stanford University.
- Büchner, M. and B. T. Kelly. (2022). “A factor model for option returns”. *Journal of Financial Economics*.
- Bybee, L., L. Gomes, and J. Valente. (2023a). “Macro-based factors for the cross-section of currency returns”.
- Bybee, L., B. T. Kelly, A. Manela, and D. Xiu. (2020). “The structure of economic news”. *Tech. rep.* National Bureau of Economic Research.
- Bybee, L., B. T. Kelly, and Y. Su. (2023b). “Narrative asset pricing: Interpretable systematic risk factors from news text”. *Review of Financial Studies*.
- Cai, T. and W. Liu. (2011). “Adaptive Thresholding for Sparse Covariance Matrix Estimation”. *Journal of the American Statistical Association*. 106: 672–684.
- Campbell, J. Y. and J. H. Cochrane. (1999). “By force of habit: A consumption-based explanation of aggregate stock market behavior”. *Journal of political Economy*. 107(2): 205–251.

- Campbell, J. Y. and S. B. Thompson. (2008). “Predicting excess stock returns out of sample: Can anything beat the historical average?” *The Review of Financial Studies*. 21(4): 1509–1531.
- Campbell, J. Y. and R. J. Shiller. (1988). “Stock Prices, Earnings, and Expected Dividends”. *The Journal of Finance*. 43(3): 661–676.
- Cao, Y., X. Liu, and J. Zhai. (2021). “Option valuation under no-arbitrage constraints with neural networks”. *European Journal of Operational Research*. 293(1): 361–374.
- Cenesizoglu, T. and A. Timmermann. (2012). “Do return prediction models add economic value?” *Journal of Banking & Finance*. 36(11): 2974–2987.
- Chamberlain, G. and M. Rothschild. (1983). “Arbitrage, Factor Structure, and Mean-Variance Analysis on Large Asset Markets”. *Econometrica*. 51: 1281–1304.
- Chatelais, N., A. Stalla-Bourdillon, and M. D. Chinn. (2023). “Forecasting real activity using cross-sectoral stock market information”. *Journal of International Money and Finance*. 131: 102800.
- Chen, A. Y. and T. Zimmermann. (2022). “Open Source Cross-Sectional Asset Pricing”. *Critical Finance Review*. 11(2): 207–264.
- Chen, B., Q. Yu, and G. Zhou. (2023a). “Useful factors are fewer than you think”. Available at SSRN 3723126.
- Chen, H., W. W. Dou, and L. Kogan. (2022a). “Measuring “Dark Matter” in Asset Pricing Models”. *Journal of Finance, forthcoming*.
- Chen, J., G. Tang, J. Yao, and G. Zhou. (2022b). “Investor Attention and Stock Returns”. *Journal of Financial and Quantitative Analysis*. 57(2): 455–484.
- Chen, L., M. Pelger, and J. Zhu. (2021). “Deep learning in asset pricing”. SSRN.
- Chen, X. and S. C. Ludvigson. (2009). “Land of addicts? an empirical investigation of habit-based asset pricing models”. *Journal of Applied Econometrics*. 24(7): 1057–1093.
- Chen, Y., B. Kelly, and D. Xiu. (2023b). “Expected Returns and Large Language Models”. *Tech. rep.* University of Chicago and Yale University.

- Chernozhukov, V., D. Chetverikov, M. Demirer, E. Duflo, C. Hansen, W. K. Newey, and J. Robins. (2018). “Double/debiased Machine Learning for Treatment and Structure Parameters”. *The Econometrics Journal*. 21(1): C1–C68.
- Chib, S., L. Zhao, and G. Zhou. (2023). “Winners from winners: A tale of risk factors”. *Management Science*.
- Chinco, A., A. D. Clark-Joseph, and M. Ye. (2019). “Sparse Signals in the Cross-Section of Returns”. *Journal of Finance*. 74(1): 449–492.
- Cho, K., B. van Merriënboer, D. Bahdanau, and Y. Bengio. (2014). “On the Properties of Neural Machine Translation: Encoder–Decoder Approaches”. In: *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*. 103–111. DOI: [10.3115/v1/W14-4012](https://doi.org/10.3115/v1/W14-4012).
- Choi, D., W. Jiang, and C. Zhang. (2022). “Alpha Go Everywhere: Machine Learning and International Stock Returns”. *Tech. rep.* The Chinese University of Hong Kong.
- Chong, E., C. Han, and F. C. Park. (2017). “Deep learning networks for stock market analysis and prediction: Methodology, data representations, and case studies”. *Expert Systems with Applications*. 83: 187–205.
- Cochrane, J. H. (2009). *Asset pricing: Revised edition*. Princeton university press.
- Cochrane, J. H. (2008). “The Dog That Did Not Bark: A Defense of Return Predictability”. *The Review of Financial Studies*. 21(4): 1533–1575.
- Cochrane, J. H. and M. Piazzesi. (2005). “Bond Risk Premia”. *American Economic Review*. 95(1): 138–160.
- Cong, L. W., G. Feng, J. He, and X. He. (2022). “Asset Pricing with Panel Tree Under Global Split Criteria”. *Tech. rep.* City University of Hong Kong.
- Cong, L. W., K. Tang, J. Wang, and Y. Zhang. (2020). “AlphaPortfolio for Investment and Economically Interpretable AI”. *Available at SSRN*.
- Connor, G., M. Hagmann, and O. Linton. (2012). “Efficient semiparametric estimation of the Fama–French model and extensions”. *Econometrica*. 80(2): 713–754.

- Connor, G. and R. A. Korajczyk. (1986). “Performance Measurement with the Arbitrage Pricing Theory: A New Framework for Analysis”. *Journal of Financial Economics*. 15(3): 373–394.
- Connor, G. and R. A. Korajczyk. (1988). “Risk and return in an equilibrium APT: Application of a new test methodology”. *Journal of financial economics*. 21(2): 255–289.
- Correia, M., J. Kang, and S. Richardson. (2018). “Asset volatility”. *Review of Accounting Studies*. 23(1): 37–94.
- Corsi, F. (2009). “A simple approximate long-memory model of realized volatility”. *Journal of Financial Econometrics*. 7: 174–196.
- Cowles, A. 3. (1933). “Can Stock Market Forecasters Forecast?” *Econometrica*. 1(3): 309–324.
- Cujean, J. and M. Hasler. (2017). “Why Does Return Predictability Concentrate in Bad Times?” *The Journal of Finance*. 72(6): 2717–2758.
- Cybenko, G. (1989). “Approximation by superpositions of a sigmoidal function”. *Mathematics of control, signals and systems*. 2(4): 303–314.
- Da, R., S. Nagel, and D. Xiu. (2022). “The Statistical Limit of Arbitrage”. *Tech. rep.* Chicago Booth.
- Das, S. R. *et al.* (2014). “Text and context: Language analytics in finance”. *Foundations and Trends® in Finance*. 8(3): 145–261.
- Davis, S. J., S. Hansen, and C. Seminario-Amez. (2020). “Firm-Level Risk Exposures and Stock Returns in the Wake of COVID-19”. *Tech. rep.* University of Chicago.
- DeMiguel, V., A. Martin-Utrera, F. J. Nogales, and R. Uppal. (2020). “A Transaction-Cost Perspective on the Multitude of Firm Characteristics”. *The Review of Financial Studies*. 33(5): 2180–2222.
- Deng, W., L. Gao, B. Hu, and G. Zhou. (2022). “Seeing is Believing: Annual Report”. *Available at SSRN 3723126*.
- Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova. (2018). “Bert: Pre-training of deep bidirectional transformers for language understanding”. *arXiv preprint arXiv:1810.04805*.
- Didisheim, A., S. Ke, B. Kelly, and S. Malamud. (2023). “Complexity in Factor Pricing Models”. *Tech. rep.* Yale University.



- Easley, D., M. López de Prado, M. O'Hara, and Z. Zhang. (2020). "Microstructure in the Machine Age". *The Review of Financial Studies*. 34(7): 3316–3363.
- Elliott, G., A. Gargano, and A. Timmermann. (2013). "Complete subset regressions". *Journal of Econometrics*. 177(2): 357–373.
- Erel, I., L. H. Stern, C. Tan, and M. S. Weisbach. (2021). "Selecting directors using machine learning". *The Review of Financial Studies*. 34(7): 3226–3264.
- Fabozzi, F. J., D. Huang, and G. Zhou. (2010). "Robust portfolios: contributions from operations research and finance". *Annals of Operations Research*. 176(1): 191–220.
- Fama, E. F. and K. R. French. (1993). "Common risk factors in the returns on stocks and bonds". *Journal of financial economics*. 33(1): 3–56.
- Fama, E. F. and K. R. French. (2010). "Luck versus skill in the cross-section of mutual fund returns". *The Journal of Finance*. 65(5): 1915–1947.
- Fama, E. F. and K. R. French. (2015). "A five-factor asset pricing model". *Journal of financial economics*. 116(1): 1–22.
- Fama, E. F. (1970). "Efficient Capital Markets: A Review of Theory and Empirical Work". *The Journal of Finance*. 25(2): 383–417.
- Fama, E. F. (1990). "Stock Returns, Expected Returns, and Real Activity". *The Journal of Finance*. 45(4): 1089–1108.
- Fama, E. F. and R. R. Bliss. (1987). "The Information in Long-Maturity Forward Rates". *The American Economic Review*. 77(4): 680–692.
- Fama, E. F. and K. R. French. (1992). "The Cross-Section of Expected Stock Returns". *The Journal of Finance*. 47: 427–465.
- Fama, E. F. and J. D. Macbeth. (1973). "Risk, Return, and Equilibrium: Empirical Tests". *Journal of Political Economy*. 81(3): 607–636.
- Fan, J., Y. Fan, and J. Lv. (2008). "High Dimensional Covariance Matrix Estimation using a Factor Model". *Journal of Econometrics*. 147: 186–197.
- Fan, J., A. Furger, and D. Xiu. (2016a). "Incorporating Global Industrial Classification Standard into Portfolio Allocation: A Simple Factor-Based Large Covariance Matrix Estimator with High Frequency Data". *Journal of Business and Economic Statistics*. 34(4): 489–503.

- Fan, J., Y. Liao, and M. Mincheva. (2013). “Large Covariance Estimation by Thresholding Principal Orthogonal Complements”. *Journal of the Royal Statistical Society, B*. 75: 603–680.
- Fan, J., Y. Liao, and W. Wang. (2016b). “Projected principal component analysis in factor models”. *Annals of Statistics*. 44(1): 219.
- Fan, J., Y. Liao, and J. Yao. (2015). “Power Enhancement in High-Dimensional Cross-Sectional Tests”. *Econometrica*. 83(4): 14977–1541.
- Feng, G., S. Giglio, and D. Xiu. (2020). “Taming the Factor Zoo: A Test of New Factors”. *Journal of Finance*. 75(3): 1327–1370.
- Feng, G., J. He, and N. G. Polson. (2018). “Deep learning for predicting asset returns”. *arXiv preprint arXiv:1804.09314*.
- Freyberger, J., A. Neuhierl, and M. Weber. (2020). “Dissecting characteristics nonparametrically”. *The Review of Financial Studies*. 33(5): 2326–2377.
- Friedman, J. H. (2001). “Greedy function approximation: a gradient boosting machine”. *Annals of statistics*: 1189–1232.
- Frost, P. A. and J. E. Savarino. (1986). “An Empirical Bayes Approach to Efficient Portfolio Selection”. *The Journal of Financial and Quantitative Analysis*. 21(3): 293–305. (Accessed on 01/30/2023).
- Fuster, A., P. Goldsmith-Pinkham, T. Ramadorai, and A. Walther. (2022). “Predictably unequal? The effects of machine learning on credit markets”. *The Journal of Finance*. 77(1): 5–47.
- Gabaix, X. (2012). “Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance”. *The Quarterly Journal of Economics*. 127: 645–700.
- Gagliardini, P., E. Ossola, and O. Scaillet. (2016). “Time-varying risk premium in large cross-sectional equity data sets”. *Econometrica*. 84(3): 985–1046.
- Gagliardini, P. and D. Ronchetti. (2019). “Comparing Asset Pricing Models by the Conditional Hansen-Jagannathan Distance\*”. *Journal of Financial Econometrics*. 18(2): 333–394.
- Garcia, D., X. Hu, and M. Rohrer. (2022). “The colour of finance words”. *Tech. rep.* University of Colorado at Boulder.

- Garleanu, N. and L. H. Pedersen. (2013). “Dynamic Trading with Predictable Returns and Transaction Costs”. *The Journal of Finance*. 68(6): 2309–2340.
- Gentzkow, M., B. Kelly, and M. Taddy. (2019). “Text as data”. *Journal of Economic Literature*. 57(3): 535–74.
- Geweke, J. and G. Zhou. (1996). “Measuring the pricing error of the arbitrage pricing theory”. *The review of financial studies*. 9(2): 557–587.
- Gibbons, M. R., S. A. Ross, and J. Shanken. (1989). “A test of the efficiency of a given portfolio”. *Econometrica: Journal of the Econometric Society*: 1121–1152.
- Giglio, S., B. Kelly, and D. Xiu. (2022a). “Factor Models, Machine Learning, and Asset Pricing”. *Annual Review of Financial Economics*. 14: 1–32.
- Giglio, S., Y. Liao, and D. Xiu. (2021a). “Thousands of Alpha Tests”. *Review of Financial Studies*. 34(7): 3456–3496.
- Giglio, S. and D. Xiu. (2021). “Asset Pricing with Omitted Factors”. *Journal of Political Economy*. 129(7): 1947–1990.
- Giglio, S., D. Xiu, and D. Zhang. (2021b). “Test Assets and Weak Factors”. *Tech. rep.* Yale University and University of Chicago.
- Giglio, S., D. Xiu, and D. Zhang. (2022b). “Prediction when Factors are Weak”. *Tech. rep.* Yale University and University of Chicago.
- Glaeser, E. L., M. S. Kincaid, and N. Naik. (2018). “Computer Vision and Real Estate: Do Looks Matter and Do Incentives Determine Looks”. *Tech. rep.* Harvard University.
- Goodfellow, I., Y. Bengio, and A. Courville. (2016). *Deep learning*. MIT press.
- Goulet Coulombe, P. and M. Göbel. (2023). “Maximally Machine-Learnable Portfolios”. *Available at SSRN 4428178*.
- Goyal, A. and A. Saretto. (2022). “Are Equity Option Returns Abnormal? IPCA Says No”. *IPCA Says No (August 19, 2022)*.
- Gu, S., B. Kelly, and D. Xiu. (2020a). “Autoencoder Asset Pricing Models”. *Journal of Econometrics*.
- Gu, S., B. Kelly, and D. Xiu. (2020b). “Empirical asset pricing via machine learning”. *The Review of Financial Studies*. 33(5): 2223–2273.

- Guijarro-Ordóñez, J., M. Pelger, and G. Zanolli. (2022). “Deep Learning Statistical Arbitrage”. *Tech. rep.* Stanford University.
- Hambly, B., R. Xu, and H. Yang. (2022). “Recent Advances in Reinforcement Learning in Finance”. *Tech. rep.* University of Oxford.
- Hansen, L. P. and R. Jagannathan. (1997). “Assessing Specification Errors in Stochastic Discount Factor Models”. *Journal of Finance*. 52: 557–590.
- Hansen, L. P. and S. F. Richard. (1987). “The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models”. *Econometrica: Journal of the Econometric Society*: 587–613.
- Hansen, L. P. and K. J. Singleton. (1982). “Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models”. *Econometrica*. 50(5): 1269–1286.
- Hansen, P. R. and A. Timmermann. (2012). “Choice of sample split in out-of-sample forecast evaluation”.
- Harvey, C. R. and Y. Liu. (2020). “False (and missed) discoveries in financial economics”. *Journal of Finance*. 75(5): 2503–2553.
- Harvey, C. R., Y. Liu, and H. Zhu. (2016). “... And the cross-section of expected returns”. *Review of Financial Studies*. 29(1): 5–68.
- Harvey, C. R. (2017). “Presidential Address: The Scientific Outlook in Financial Economics”. *Journal of Finance*. 72(4): 1399–1440.
- Harvey, C. R. and W. E. Ferson. (1999). “Conditioning Variables and the Cross-Section of Stock Returns”. *Journal of Finance*. 54: 1325–1360.
- Hastie, T., A. Montanari, S. Rosset, and R. J. Tibshirani. (2019). “Surprises in high-dimensional ridgeless least squares interpolation”. *arXiv preprint arXiv:1903.08560*.
- Haugen, R. A. and N. L. Baker. (1996). “Commonality in the determinants of expected stock returns”. *Journal of Financial Economics*. 41(3): 401–439.
- Hayek, F. A. (1945). “The Use of Knowledge in Society”. *The American Economic Review*. 35(4): 519–530.
- He, A., S. He, D. Rapach, and G. Zhou. (2022a). “Expected Stock Returns in the Cross-section: An Ensemble Approach”. *Working Paper*.

- He, K., X. Zhang, S. Ren, and J. Sun. (2016). “Deep Residual Learning for Image Recognition”. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 770–778. DOI: [10.1109/CVPR.2016.90](https://doi.org/10.1109/CVPR.2016.90).
- He, S., M. Yuan, and G. Zhou. (2022b). “Principal Portfolios: A Note”. *Working Paper*.
- He, X., G. Feng, J. Wang, and C. Wu. (2021). “Predicting Individual Corporate Bond Returns”. *Tech. rep.* City University of Hong Kong.
- He, Z., B. Kelly, and A. Manela. (2017). “Intermediary asset pricing: New evidence from many asset classes”. *Journal of Financial Economics*. 126(1): 1–35.
- Hochreiter, S. and J. Schmidhuber. (1997). “Long short-term memory”. *Neural Computation*. 9: 1735–1780.
- Hong, H., T. Lim, and J. C. Stein. (2000). “Bad News Travels Slowly: Size, Analyst Coverage, and the Profitability of Momentum Strategies”. *The Journal of Finance*. 55(1): 265–295.
- Hornik, K., M. Stinchcombe, and H. White. (1989). “Multilayer feedforward networks are universal approximators”. *Neural networks*. 2(5): 359–366.
- Hornik, K., M. Stinchcombe, and H. White. (1990). “Universal approximation of an unknown mapping and its derivatives using multilayer feedforward networks”. *Neural networks*. 3(5): 551–560.
- Hou, K., C. Xue, and L. Zhang. (2018). “Replicating Anomalies”. *The Review of Financial Studies*. 33(5): 2019–2133.
- Hu, A. and S. Ma. (2020). “Human interactions and financial investment: A video-based approach”. *Available at SSRN*.
- Huang, D., F. Jiang, K. Li, G. Tong, and G. Zhou. (2022). “Scaled PCA: A New Approach to Dimension Reduction”. *Management Science*. 68(3): 1591–2376.
- Huang, D., F. Jiang, J. Tu, and G. Zhou. (2014). “Investor Sentiment Aligned: A Powerful Predictor of Stock Returns”. *The Review of Financial Studies*. 28(3): 791–837.
- Huang, J., J. L. Horowitz, and F. Wei. (2010). “Variable selection in nonparametric additive models”. *The Annals of Statistics*. 38(4): 2282–2313.

- Huberman, G. (1982). “A Simple Approach to Arbitrage Pricing Theory”. *Journal of Economic Theory*. 28(1): 183–191.
- Ingersoll, J. E. (1984). “Some Results in the Theory of Arbitrage Pricing”. *Journal of Finance*. 39(4): 1021–1039.
- Israel, R., B. Kelly, and T. J. Moskowitz. (2020). “Can Machines “Learn” Finance?” *Journal of Investment Management*. 18(2): 23–36.
- Israelov, R. and B. T. Kelly. (2017). “Forecasting the distribution of option returns”. Available at SSRN 3033242.
- Jacot, A., F. Gabriel, and C. Hongler. (2018). “Neural tangent kernel: Convergence and generalization in neural networks”. *arXiv preprint arXiv:1806.07572*.
- Jegadeesh, N. and D. Wu. (2013). “Word power: A new approach for content analysis”. *Journal of Financial Economics*. 110(3): 712–729.
- Jensen, T. I., B. Kelly, C. Seminario-Amez, and L. H. Pedersen. (2022). “Machine Learning and the Implementable Efficient Frontier”. *Tech. rep.* Copenhagen Business School.
- Jensen, T. I., B. Kelly, and L. H. Pedersen. (2023). “Is There a Replication Crisis in Finance?” *Journal of Finance*. 78(5): 2465–2518.
- Jiang, F., G. Tang, and G. Zhou. (2018). “Firm characteristics and Chinese stocks”. *Journal of Management Science and Engineering*. 3(4): 259–283.
- Jiang, J., B. Kelly, and D. Xiu. (2022). “(Re-)Imag(in)ing Price Trends”. *Journal of Finance, forthcoming*.
- Jobson, J. D. and B. Korkie. (1980). “Estimation for Markowitz Efficient Portfolios”. *Journal of the American Statistical Association*. 75(371): 544–554. (Accessed on 01/30/2023).
- Jorion, P. (1986). “Bayes-Stein Estimation for Portfolio Analysis”. *The Journal of Financial and Quantitative Analysis*. 21(3): 279–292. (Accessed on 01/30/2023).
- Jurado, K., S. C. Ludvigson, and S. Ng. (2015). “Measuring uncertainty”. *The American Economic Review*. 105(3): 1177–1216.
- Kan, R., X. Wang, and G. Zhou. (2022). “Optimal Portfolio Choice with Estimation Risk: No Risk-free Asset Case”. *Management Science*. 68(3): 1591–2376.
- Kan, R. and C. Zhang. (1999). “Two-Pass Tests of Asset Pricing Models with Useless Factors”. *The Journal of Finance*. 54(1): 203–235.

- Kan, R. and G. Zhou. (2007). “Optimal Portfolio Choice with Parameter Uncertainty”. *Journal of Financial and Quantitative Analysis*. 42(3): 621–656.
- Ke, T., B. Kelly, and D. Xiu. (2019). “Predicting Returns with Text Data”. *Tech. rep.* Harvard University, Yale University, and the University of Chicago.
- Kelly, B., S. Malamud, and L. H. Pedersen. (2020a). “Principal Portfolios”. *Working Paper*.
- Kelly, B., S. Malamud, and K. Zhou. (2022a). “Virtue of Complexity in Return Prediction”. *Tech. rep.* Yale University.
- Kelly, B., A. Manela, and A. Moreira. (2018). “Text Selection”. *Working paper*.
- Kelly, B., T. Moskowitz, and S. Pruitt. (2021). “Understanding Momentum and Reversal”. *Journal of Financial Economics*. 140(3): 726–743.
- Kelly, B., D. Palhares, and S. Pruitt. (2023). “Modeling Corporate Bond Returns”. *Journal of Finance*. 78(4): 1967–2008.
- Kelly, B. and S. Pruitt. (2013). “Market expectations in the cross-section of present values”. *The Journal of Finance*. 68(5): 1721–1756.
- Kelly, B. and S. Pruitt. (2015). “The three-pass regression filter: A new approach to forecasting using many predictors”. *Journal of Econometrics*. 186(2): 294–316.
- Kelly, B., S. Pruitt, and Y. Su. (2020b). “Characteristics are Covariances: A Unified Model of Risk and Return”. *Journal of Financial Economics*.
- Kelly, B. T., S. Malamud, and K. Zhou. (2022b). “The Virtue of Complexity Everywhere”. *Available at SSRN*.
- Kim, S., R. Korajczyk, and A. Neuhierl. (2021). “Arbitrage Portfolios”. *Review of Financial Studies*. 34(6): 2813–2856.
- Koopmans, T. C. (1947). “Measurement without theory”. *The Review of Economics and Statistics*. 29(3): 161–172.
- Kosowski, R., A. Timmermann, R. Wermers, and H. White. (2006). “Can mutual fund “stars” really pick stocks? New evidence from a bootstrap analysis”. *The Journal of Finance*. 61(6): 2551–2595.
- Kozak, S. (2020). “Kernel trick for the cross-section”. *Available at SSRN 3307895*.

- Kozak, S., S. Nagel, and S. Santosh. (2018). “Interpreting factor models”. *The Journal of Finance*. 73(3): 1183–1223.
- Kozak, S., S. Nagel, and S. Santosh. (2020). “Shrinking the cross-section”. *Journal of Financial Economics*. 135(2): 271–292.
- Langlois, H. (2023). “What matters in a characteristic?” *Journal of Financial Economics*. 149(1): 52–72.
- Ledoit, O. and M. Wolf. (2004). “Honey, I shrunk the sample covariance matrix”. *Journal of Portfolio Management*. 30: 110–119.
- Ledoit, O. and M. Wolf. (2012). “Nonlinear shrinkage estimation of large-dimensional covariance matrices”. *The Annals of Statistics*. 40: 1024–1060.
- Leippold, M., Q. Wang, and W. Zhou. (2022). “Machine learning in the Chinese stock market”. *Journal of Financial Economics*. 145(2, Part A): 64–82.
- Lettau, M. and S. Ludvigson. (2001). “Consumption, Aggregate Wealth, and Expected Stock Returns”. *The Journal of Finance*. 56(3): 815–849.
- Lettau, M. and M. Pelger. (2020a). “Estimating Latent Asset-Pricing Factors”. *Journal of Econometrics*. 218: 1–31.
- Lettau, M. and M. Pelger. (2020b). “Factors that fit the time series and cross-section of stock returns”. *Review of Financial Studies*. 33(5): 2274–2325.
- Lewellen, J. (2015). “The Cross-section of Expected Stock Returns”. *Critical Finance Review*. 4(1): 1–44.
- Li, K., F. Mai, R. Shen, and X. Yan. (2021). “Measuring corporate culture using machine learning”. *The Review of Financial Studies*. 34(7): 3265–3315.
- Li, S. Z. and Y. Tang. (2022). “Automated Risk Forecasting”. *Tech. rep.* Rutgers, The State University of New Jersey.
- Light, N., D. Maslov, and O. Rytchkov. (2017). “Aggregation of Information About the Cross Section of Stock Returns: A Latent Variable Approach”. *The Review of Financial Studies*. 30(4): 1339–1381.
- Lo, A. W. and A. C. MacKinlay. (1990). “Data-snooping biases in tests of financial asset pricing models”. *Review of financial studies*. 3(3): 431–467.



- Loughran, T. and B. McDonald. (2011). “When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks”. *The Journal of Finance*. 66(1): 35–65.
- Loughran, T. and B. McDonald. (2020). “Textual analysis in finance”. *Annual Review of Financial Economics*. 12: 357–375.
- Lucas Jr, R. E. (1976). “Econometric policy evaluation: A critique”. In: *Carnegie-Rochester conference series on public policy*. Vol. 1. North-Holland. 19–46.
- Ludvigson, S. C. and S. Ng. (2010). “A factor analysis of bond risk premia”. In: *Handbook of empirical economics and finance*. Ed. by A. Ulah and D. E. A. Giles. Vol. 1. Chapman and Hall, Boca Raton, FL. Chap. 12. 313–372.
- Ludvigson, S. C. (2013). “Chapter 12 - Advances in Consumption-Based Asset Pricing: Empirical Tests”. In: ed. by G. M. Constantinides, M. Harris, and R. M. Stulz. Vol. 2. *Handbook of the Economics of Finance*. Elsevier. 799–906.
- Ludvigson, S. C. and S. Ng. (2007). “The empirical risk–return relation: A factor analysis approach”. *Journal of Financial Economics*. 83(1): 171–222.
- Lynch, A. W. and P. Balduzzi. (2000). “Predictability and Transaction Costs: The Impact on Rebalancing Rules and Behavior”. *The Journal of Finance*. 55(5): 2285–2309.
- Lyonnet, V. and L. H. Stern. (2022). “Venture Capital (Mis) allocation in the Age of AI”. Available at SSRN 4260882.
- Malloy, C. J., T. J. Moskowitz, and A. Vissing-Jorgensen. (2009). “Long-run stockholder consumption risk and asset returns”. *The Journal of Finance*. 64(6): 2427–2479.
- Manela, A. and A. Moreira. (2017). “News implied volatility and disaster concerns”. *Journal of Financial Economics*. 123(1): 137–162.
- Markowitz, H. (1952). “Portfolio selection”. *Journal of Finance*. 7(1): 77–91.
- Martin, I. W. and S. Nagel. (2021). “Market efficiency in the age of big data”. *Journal of Financial Economics*.
- Mehra, R. and E. C. Prescott. (1985). “The equity premium: A puzzle”. *Journal of Monetary Economics*. 15(2): 145–161.

- Menzly, L., T. Santos, and P. Veronesi. (2004). “Understanding Predictability”. *Journal of Political Economy*. 112(1): 1–47. (Accessed on 02/06/2023).
- Merton, R. C. (1973). “An Intertemporal Capital Asset Pricing Model”. *Econometrica*. 41: 867–887.
- Michaud, R. O. (1989). “The Markowitz Optimization Enigma: Is ‘Optimized’ Optimal?” *Financial Analysts Journal*. 45(1): 31–42. (Accessed on 01/30/2023).
- Mittnik, S., N. Robinzonov, and M. Spindler. (2015). “Stock market volatility: Identifying major drivers and the nature of their impact”. *Journal of Banking & Finance*. 58: 1–14.
- Moritz, B. and T. Zimmermann. (2016). “Tree-Based Conditional Portfolio Sorts: The Relation Between Past and Future Stock Returns”. *Tech. rep.* Ludwig Maximilian University Munich.
- Nagel, S. and K. Singleton. (2011). “Estimation and Evaluation of Conditional Asset Pricing Models”. *The Journal of Finance*. 66(3): 873–909. (Accessed on 02/20/2023).
- Nishii, R. (1984). “Asymptotic Properties of Criteria for Selection of Variables in Multiple Regression”. *The Annals of Statistics*. 12(2): 758–765.
- Novy-Marx, R. (2014). “Predicting anomaly performance with politics, the weather, global warming, sunspots, and the stars”. *Journal of Financial Economics*. 112(2): 137–146.
- Obaid, K. and K. Pukthuanthong. (2022). “A picture is worth a thousand words: Measuring investor sentiment by combining machine learning and photos from news”. *Journal of Financial Economics*. 144: 273–297.
- Ohlson, J. A. (1980). “Financial Ratios and the Probabilistic Prediction of Bankruptcy”. *Journal of Accounting Research*. 18(1): 109–131.
- Onatski, A. (2009). “Testing hypotheses about the number of factors in large factor models”. *Econometrica*. 77(5): 1447–1479.
- Onatski, A. (2010). “Determining the Number of Factors from Empirical Distribution of Eigenvalues”. *Review of Economics and Statistics*. 92: 1004–1016.

- Onatski, A. (2012). “Asymptotics of the principal components estimator of large factor models with weakly influential factors”. *Journal of Econometrics*. 168: 244–258.
- Pastor, L. (2000). “Portfolio Selection and Asset Pricing Models”. *The Journal of Finance*. 55(1): 179–223.
- Pastor, L. and R. F. Stambaugh. (2000). “Comparing Asset Pricing Models: An Investment Perspective”. *Journal of Financial Economics*. 56: 335–381.
- Pástor, L. and R. F. Stambaugh. (2003). “Liquidity Risk and Expected Stock Returns”. *Journal of Political Economy*. 111(3): 642–685.
- Pesaran, H. and T. Yamagata. (2017). “Testing for Alpha in Linear Factor Pricing Models with a Large Number of Securities”. *Tech. rep.*
- Petersen, M. A. (2008). “Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches”. *The Review of Financial Studies*. 22(1): 435–480.
- Pettenuzzo, D., A. Timmermann, and R. Valkanov. (2014). “Forecasting stock returns under economic constraints”. *Journal of Financial Economics*. 114(3): 517–553.
- Pukthuanthong, K., R. Roll, and A. Subrahmanyam. (2019). “A Protocol for Factor Identification”. *Review of Financial Studies*. 32(4): 1573–1607.
- Radford, A., J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, *et al.* (2019). “Language models are unsupervised multitask learners”. *OpenAI blog*. 1(8): 9.
- Rahimi, A. and B. Recht. (2007). “Random Features for Large-Scale Kernel Machines.” In: *NIPS*. Vol. 3. No. 4. Citeseer. 5.
- Rapach, D. and G. Zhou. (2013). “Chapter 6 - Forecasting Stock Returns”. In: *Handbook of Economic Forecasting*. Ed. by G. Elliott and A. Timmermann. Vol. 2. *Handbook of Economic Forecasting*. Elsevier. 328–383.
- Rapach, D. and G. Zhou. (2022). “Asset pricing: Time-series predictability”.
- Rapach, D. E., J. K. Strauss, and G. Zhou. (2010). “Out-of-sample equity premium prediction: Combination forecasts and links to the real economy”. *The Review of Financial Studies*. 23(2): 821–862.

- Rapach, D. E., J. K. Strauss, and G. Zhou. (2013). "International stock return predictability: what is the role of the United States?" *The Journal of Finance*. 68(4): 1633–1662.
- Rather, A. M., A. Agarwal, and V. Sastry. (2015). "Recurrent neural network and a hybrid model for prediction of stock returns". *Expert Systems with Applications*. 42(6): 3234–3241.
- Roll, R. (1977). "A Critique of the Asset Pricing Theory's Tests". *Journal of Financial Economics*. 4: 129–176.
- Rosenberg, B. (1974). "Extra-Market Components of Covariance in Security Returns". *Journal of Financial and Quantitative Analysis*. 9(2): 263–274.
- Rosenberg, J. V. and R. F. Engle. (2002). "Empirical pricing kernels". *Journal of Financial Economics*. 64(3): 341–372.
- Ross, S. A. (1976). "The Arbitrage Theory of Capital Asset Pricing". *Journal of Economic Theory*. 13(3): 341–360.
- Rossi, A. G. (2018). "Predicting stock market returns with machine learning". *Georgetown University*.
- Rossi, A. G. and A. Timmermann. (2015). "Modeling Covariance Risk in Merton's ICAPM". *The Review of Financial Studies*. 28(5): 1428–1461.
- Samuelson, P. A. (1965). "Rational Theory of Warrant Pricing". *Industrial Management Review*. 6(2): 13–39.
- Schaller, H. and S. V. Norden. (1997). "Regime switching in stock market returns". *Applied Financial Economics*. 7(2): 177–191.
- Schapire, R. E. (1990). "The Strength of Weak Learnability". *Machine Learning*. 5(2): 197–227.
- Sezer, O. B., M. U. Gudelek, and A. M. Ozbayoglu. (2020). "Financial time series forecasting with deep learning : A systematic literature review: 2005–2019". *Applied Soft Computing*. 90: 106–181.
- Shanken, J. (1992a). "On the Estimation of Beta Pricing Models". *Review of Financial Studies*. 5: 1–33.
- Shanken, J. (1992b). "The Current State of the Arbitrage Pricing Theory". *Journal of Finance*. 47(4): 1569–1574.
- Shiller, R. J. (1981). "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" *The American Economic Review*. 71(3): 421–436.

- Simon, F., S. Weibels, and T. Zimmermann. (2022). “Deep Parametric Portfolio Policies”. *Tech. rep.* University of Cologne.
- Singh, R. and S. Srivastava. (2017). “Stock prediction using deep learning”. *Multimedia Tools and Applications*. 76(18): 18569–18584.
- Spigler, S., M. Geiger, S. d’Ascoli, L. Sagun, G. Biroli, and M. Wyart. (2019). “A jamming transition from under-to over-parametrization affects generalization in deep learning”. *Journal of Physics A: Mathematical and Theoretical*. 52(47): 474001.
- Stock, J. H. and M. W. Watson. (2002). “Forecasting Using Principal Components from a Large Number of Predictors”. *Journal of the American Statistical Association*. 97(460): 1167–1179.
- Stone, M. (1977). “An Asymptotic Equivalence of Choice of Model by Cross-Validation and Akaike’s Criterion”. *Journal of the Royal Statistical Society Series B*. 39(1): 44–47.
- Sullivan, R., A. Timmermann, and H. White. (1999). “Data-snooping, technical trading rule performance, and the bootstrap”. *The Journal of Finance*. 54(5): 1647–1691.
- Taddy, M. (2013). “Multinomial inverse regression for text analysis”. *Journal of American Statistical Association*. 108(503): 755–770.
- Tetlock, P. C. (2007). “Giving Content to Investor Sentiment: The Role of Media in the Stock Market”. *The Journal of Finance*. 62(3): 1139–1168.
- Tu, J. and G. Zhou. (2010). “Incorporating Economic Objectives into Bayesian Priors: Portfolio Choice under Parameter Uncertainty”. *The Journal of Financial and Quantitative Analysis*. 45(4): 959–986. (Accessed on 01/30/2023).
- Van Binsbergen, J. H. and R. Koijen. (2010). “Predictive Regressions: A Present-Value Approach”. *The Journal of Finance*. 65(4): 1439–1471.
- Van Binsbergen, J. H. and C. C. Opp. (2019). “Real anomalies”. *Journal of Finance*. 74(4): 1659–1706.
- Wachter, J. (2013). “Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?” *The Journal of Finance*. 68: 987–1035.

- Wachter, J. A. (2006). “A consumption-based model of the term structure of interest rates”. *Journal of Financial Economics*. 79(2): 365–399.
- Welch, I. and A. Goyal. (2008). “A comprehensive look at the empirical performance of equity premium prediction”. *The Review of Financial Studies*. 21(4): 1455–1508.
- White, H. (2000). “A reality check for data snooping”. *Econometrica*. 68(5): 1097–1126.
- Windmüller, S. (2022). “Firm characteristics and global stock returns: A conditional asset pricing model”. *The Review of Asset Pricing Studies*. 12(2): 447–499.
- Yogo, M. (2006). “A Consumption-Based Explanation of Expected Stock Returns”. *The Journal of Finance*. 61(2): 539–580.
- Yuan, M. and G. Zhou. (2022). “Why Naive 1/N Diversification Is Not So Naive, and How to Beat It?” Available at SSRN.
- Zhang, S., S. Roller, N. Goyal, M. Artetxe, M. Chen, S. Chen, C. Dewan, M. Diab, X. Li, X. V. Lin, T. Mihaylov, M. Ott, S. Shleifer, K. Shuster, D. Simig, P. S. Koura, A. Sridhar, T. Wang, and L. Zettlemoyer. (2022). “OPT: Open Pre-trained Transformer Language Models”. arXiv: [2205.01068](https://arxiv.org/abs/2205.01068) [cs.CL].